

11-13-08

A

11/10/00
JC715 U.S. PTO

Please type a plus sign (+) inside this box [+]

PTO/SB/05 (11-00)

Approved for use through 10/31/2002. OMB 0651-0032

Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL
(Only for new nonprovisional applications under 37 CFR 1.53(b))

Attorney Docket No. 004683.P003

First Named Inventor Todd Krautkremer

Title: "APPLICATION SERVICE LEVEL MEDIATION AND METHOD OF USING THE SAME"

Express Mail Label No. EL143569335US

jc930 U.S. PTO
09/7/04
11/10/00

ADDRESS TO: Assistant Commissioner for Patents
Box Patent Application
Washington, D. C. 20231

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

1. X **Fee Transmittal Form (e.g., PTO/SB/17)**
(Submit an original, and a duplicate for fee processing)
2. **Applicant Claims Small Entity Status. (37 CFR 1.27)**
3. X **Specification (Total Pages 28)**
(preferred arrangement set forth below)
 - Descriptive Title of the Invention
 - Cross Reference to Related Applications
 - Statement Regarding Fed sponsored R & D
 - Reference sequence listing, a table,
or a computer program listing appendix
 - Background of the Invention
 - Brief Summary of the Invention
 - Brief Description of the Drawings (if filed)
 - Detailed Description
 - Claim(s)
 - Abstract of the Disclosure
4. X **Drawings(s) (35 USC 113) (Total Sheets 5)**
5. X **Oath or Declaration (Total Pages 5)**
 - a. Newly Executed (Original or Copy)
 - b. Copy from a Prior Application (37 CFR 1.63(d))
(for Continuation/Divisional with Box 17 completed)
 - i. DELETIONS OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b).
 - c. X Unsigned.
6. **Application Data Sheet. (37 CFR 1.76)**
7. CD-ROM or CD-R in duplicate, large table or Computer Program (Appendix)
8. Nucleotide and/or Amino Acid Sequence Submission
(if applicable, all necessary)
 - a. Computer Readable Form (CRF)
 - b. Specification Sequence Listing on:
 - i. CD-ROM or CD-R (2 copies); or
 - ii. paper
 - c. Statement verifying identity of above copies

11/08/00

ACCOMPANYING APPLICATION PARTS

9. ☐ **Assignment Papers (cover sheet & documents(s))**
10. ☐ a. Separate 37 CFR 3.73(b) Statement (where there is an assignee)
- ☐ b. Power of Attorney
11. ☐ English Translation Document (if applicable)
12. ☐ a. Information Disclosure Statement (IDS)/PTO-1449
- ☐ b. Copies of IDS Citations
13. ☐ **Preliminary Amendment**
14. ☒ **Return Receipt Postcard (MPEP 503) (Should be specifically itemized)**
15. ☐ Certified Copy of Priority Document(s) (if foreign priority is claimed)
16. ☐ Request and Certification under 35 U.S.C. 122(b)(2)(B)(i). Applicant must attach form PTO/SB/35 or its equivalent.
17. ☐ Other: _____

18A. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP)
Of Prior Application No.: _____ Examiner _____ Group Art Unit _____

(which is a ☐ continuation/ ☐ divisional/ ☐ CIP of prior application no. _____,
which is a ☐ continuation/ ☐ divisional/ ☐ CIP of prior application no. _____) (List entire chain of priority)

For CONTINUATION AND DIVISIONAL APPS only: The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 5b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

18B. Statement under 37 CFR 3.73(b) for continuing application:

The undersigned states that _____ (name of assignee) is the assignee of the entire right, title, and interest in the accompanying patent application by virtue of an assignment recorded in the Patent and Trademark Office at Reel No. _____ Frame No. _____ (or a copy of which is attached).

19. Correspondence Address

☐ Customer Number or Bar Code Label _____
or _____ (Insert Customer No. or Attach Bar Code Label here)

☒ Correspondence Address Below

NAME Michael J. Mallie, Reg No. 36,591

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

ADDRESS 12400 Wilshire Boulevard

Seventh Floor

CITY Los Angeles STATE California ZIP CODE 90025-1026

Country U.S.A. TELEPHONE (408) 720-8300 FAX (408) 720-9397

Name (PRINT/TYPE): Michael J. Mallie

Registration No.: 36,591

Signature: _____

Date: November 10, 2000

FEE TRANSMITTAL FOR FY 2001**TOTAL AMOUNT OF PAYMENT (\$)** 710.00**Complete if Known:**

Application No. Not yet assigned
 Filing Date Not yet assigned
 First Named Inventor Todd Krautkremer, et al.
 Group Art Unit Not yet assigned
 Examiner Name Not yet assigned
 Attorney Docket No. 004683.P003

METHOD OF PAYMENT (check one)

1. ☒ The Commissioner is hereby authorized to charge indicated fees and credit any over payments to:

Deposit Account Number 02-2666
 Deposit Account Name Blakely, Sokoloff, Taylor & Zafman, LLP

- ☒ Charge Any Additional Fee Required Under 37 CFR 1.16 and 1.17

- ☐ Applicant claims small entity status. See 37 CFR 1.27

2. ☒ Payment Enclosed: ☒ Check
☐ Credit Card
☐ Money Order
☐ Other

FEE CALCULATION**1. BASIC FILING FEE**

<u>Large Entity</u>		<u>Small Entity</u>		<u>Fee Description</u>	<u>Fee Paid</u>
Fee Code	Fee (\$)	Fee Code	Fee (\$)		
101	710	201	355	Utility application filing fee	<u>710.00</u>
106	320	206	160	Design application filing fee	
107	490	207	245	Plant filing fee	
108	710	208	355	Reissue filing fee	
114	150	214	75	Provisional application filing fee	

SUBTOTAL (1) \$ 710.00**2. EXTRA CLAIM FEES**

			<u>Extra Claims</u>	<u>Fee from below</u>	<u>Fee Paid</u>
Total Claims	<u>9</u>	- 20** =	<u>0</u>	X <u>18.00</u>	= <u>0</u>
Independent Claims	<u>1</u>	- 3** =	<u>0</u>	X <u>80.00</u>	= <u>0</u>
Multiple Dependent					=

****Or number previously paid, if greater; For Reissues, see below.**

<u>Large Entity</u>		<u>Small Entity</u>		<u>Fee Description</u>
Fee Code	Fee (\$)	Fee Code	Fee (\$)	
103	18	203	9	Claims in excess of 20
102	80	202	40	Independent claims in excess of 3
104	270	204	135	Multiple dependent claim, if not paid
109	80	209	40	**Reissue independent claims over original patent
110	18	210	9	**Reissue claims in excess of 20 and over original patent

SUBTOTAL (2) \$ 0

FEE CALCULATION (continued)**3. ADDITIONAL FEES**

<u>Large Entity</u>		<u>Small Entity</u>		<u>Fee Description</u>	<u>Fee Paid</u>
<u>Fee Code</u>	<u>Fee (\$)</u>	<u>Fee Code</u>	<u>Fee (\$)</u>		
105	130	205	65	Surcharge - late filing fee or oath	_____
127	50	227	25	Surcharge - late provisional filing fee or cover sheet	_____
139	130	139	130	Non-English specification	_____
147	2,520	147	2,520	For filing a request for ex parte reexamination	_____
112	920*	112	920*	Requesting publication of SIR prior to Examiner action	_____
113	1,840*	113	1,840*	Requesting publication of SIR after Examiner action	_____
115	110	215	55	Extension for reply within first month	_____
116	390	216	195	Extension for reply within second month	_____
117	890	217	445	Extension for reply within third month	_____
118	1,390	218	695	Extension for reply within fourth month	_____
128	1,890	228	945	Extension for reply within fifth month	_____
119	310	219	155	Notice of Appeal	_____
120	310	220	155	Filing a brief in support of an appeal	_____
121	270	221	135	Request for oral hearing	_____
138	1,510	138	1,510	Petition to institute a public use proceeding	_____
140	110	240	55	Petition to revive - unavoidable	_____
141	1,240	241	620	Petition to revive - unintentional	_____
142	1,240	242	620	Utility issue fee (or reissue)	_____
143	440	243	220	Design issue fee	_____
144	600	244	300	Plant issue fee	_____
122	130	122	130	Petitions to the Commissioner	_____
123	130	123	130	Petitions related to provisional applications	_____
126	180	126	180	Submission of Information Disclosure Stmt	_____
581	40	581	40	Recording each patent assignment per property (times number of properties)	_____
146	710	246	355	For filing a submission after final rejection (see 37 CFR 1.129(a))	_____
149	710	249	355	For each additional invention to be examined (see 37 CFR 1.129(b))	_____
179	710	279	355	Request for Continued Examination (RCE)	_____
169	900	169	900	Request for expedited examination of a design application	_____
Other fee (specify) _____					_____
Other fee (specify) _____					_____

SUBTOTAL (3) \$ 0

*Reduced by Basic Filing Fee Paid

SUBMITTED BY:Typed or Printed Name: Michael J. MallieSignature: _____ Date: November 10, 2000Reg. Number: 36,591 Telephone Number: 408-720-8300**WARNING:** Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

UNITED STATES PATENT APPLICATION

for

APPLICATION SERVICE LEVEL MEDIATION AND METHOD OF
USING THE SAME

Inventors:

Todd Krautkremer

Guy Riddle

prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

12400 Wilshire Boulevard

Los Angeles, CA 90025-1026

(408) 720-8598

File No.: 004683.P003

EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number: EL14356933545

Date of Deposit: 11/10/00

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Assistant Commissioner for Patents, Washington, D. C. 20231

Mara E. Brown

(Typed or printed name of person mailing paper or fee)

Mara E. Brown
(Signature of person mailing paper or fee)

11/10/00
(Date signed)

APPLICATION SERVICE LEVEL MEDIATION AND METHOD OF USING THE SAME

FIELD OF THE INVENTION

5 The present invention relates to the field of network administration; more particularly, the present invention relates to application service providers and their use of an application layer demarcation point to monitor performance.

BACKGROUND OF THE INVENTION

10 Services delivered by providers of networked application services, by their nature, span a variety of provider and customer owned and managed infrastructures. For example, the application services begin at a provider or
15 customer owned hosting platform within a provider or customer owned data center infrastructure, travel across one or more wide-area networks (WANs) owned and managed by one or more service providers and across one or more customer owned WAN and LAN infrastructures, before reaching a customer-owned desktop or mobile computing platform.

20 As application traffic traverses service boundaries between the source provider and infrastructure supplier providers and source provider and customers, the source provider's ability to effect overall service-levels is

greatly diminished. Because there is no application-aware demarcation point that aligns with the source provider's service boundary, there is no easy way to establish an actionable service-level agreement. A source provider cannot control the performance of application service over portions of

5 infrastructure it does not own or control.

Current approaches to application performance monitoring and service-level management measure end-to-end (server to desktop) performance without regard to the service provider boundaries. This lack of a clear demarcation point between provider and customer impairs the

10 provider's ability to deliver actionable service-level agreements and results in significant service costs in mediating between provider and customer for infrastructure-induced performance problems.

SUMMARY OF THE INVENTION

- A method and apparatus for using an application layer demarcation point are described. In one embodiment, the method comprises monitoring end-to-end performance of a network application
- 5 at an application demarcation point in a network, and mediating between provider infrastructure and customer infrastructure based on results of monitoring.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the invention, which, however, should not be taken
5 to limit the invention to the specific embodiments, but are for explanation and understanding only.

Figure 1 illustrates a network environment that includes an application level demarcation point.
10

Figure 2 illustrates an exemplary network environment in which an application demarcation point is utilized.

Figure 3 is an alternative view of Figure 1
15

Figure 4 illustrates one embodiment of software code executed by one embodiment of network device.

Figure 5 is a data flow diagram of one embodiment of a network
20 device.

DETAILED DESCRIPTION OF THE PRESENT INVENTION

A method and apparatus for using an application layer demarcation point are described. The application layer demarcation (demarc) point is provided to enable providers of application services to measure the

5 quantitative performance and qualitative performance (e.g., an application is too slow) of networked applications end-to-end and to mediate between performance impacts of the underlying service layers and service boundaries. Thus, a service boundary demarcation point for the application is provided that enables source providers of application services to mediate

10 between the underlying infrastructure which they own and manage (or its supplier partners) and the customer's infrastructure.

Since a service provider cannot manage customer infrastructure it does not own or control, proving an application-aware demarcation and mediation point allows a service provider to develop actionable service-level

15 agreements that align with its service boundary. In this manner, a service provider delivering service at the application layer may use a demarcation point at the application layer to enable verification that the service was delivered. This is performed without changing the clients or servers in the network.

Using the application layer demarcation allows for monitoring and measuring performance (e.g., congestion or a lack thereof) and whether performance problems (e.g., packet discard and retransmission that slows overall response time, etc.) are on the service side or on the customer side of the network. In one embodiment, in regard to congestion, identifying performance problems includes identifying the class of traffic that is being effected in order to help facilitate solving the problem for the customer and/or providing services to the customer to remove or reduce the problem.

The demarcation also allows for availability measures to be made by making determinations as, for example, to what time the network carries traffic or as to what time a host is up and running.

In one embodiment, the monitoring and mediation functionality reside on an appliance platform that is installed, usually by the service provider, at the service boundary between the providers' and customers' infrastructure. This appliance measure the end-to-end performance and service-levels of application services and determines from where a service-affecting problem is emanating. Moreover, the appliance becomes an application-aware demarcation point that allows service provider to create,

monitor and manage service-level agreements for the portion of the end-to-end service delivers that actually traverses its infrastructure.

The demarcation point allows the use of data collected relative to the demarcation point as a tool, such as a diagnostic tool to correct identified
5 problems in the network or a business tool to provide business opportunities. For example, by being able to have an application demarcation point, service providers are able to indicate that users need additional resources (e.g., bandwidth resources, additional server capacity, etc.) and are able to indicate to users in a measurable and discrete way that
10 they need to obtain those additional resources. Thus, the demarcation point allows a service provider to determine whether appropriate additional services should be presented to customers. Such services may be provided at an additional cost to the customers. Furthermore, the application demarcation point may be used to perform capacity planning, particularly
15 by being able to determine what each customer requires and ensuring that those requirements are met.

In the following description, numerous details are set forth to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art, that the present invention may be

practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form, rather than in detail, in order to avoid obscuring the present invention.

Some portions of the detailed descriptions which follow are presented

5 in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-

10 consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for

15 reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically

stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer

5 system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

10 The present invention also relates to apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable
15 storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus.

The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized
5 apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the
10 invention as described herein.

A machine-readable medium includes any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer). For example, a machine-readable medium includes read only memory ("ROM"); random access memory ("RAM"); magnetic disk storage media;
15 optical storage media; flash memory devices; electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.); etc.

Figure 1 illustrates a network environment that includes an application level demarcation point. Referring to Figure 1, an application

demarcation point 101 is located between a provider portion 110 of the network and a customer portion 111 of the network. Provider portion 110 of the network couples one or more servers (e.g., server 150) to the network, while client portion 111 of the network couples one or more clients (e.g., client 151) to the network. In one embodiment, demarcation point 101 is located logically between customer portion 111 and provider portion 110 and physically at network appliance, or device 120.

Network device 120 monitors application performance using delay metrics that characterize the delay associated with end-to-end traffic in the network. Such traffic may be that of, for example, interactive applications that are operating with requests and responses to those requests traversing the network. Apart from network delay metrics, other metrics that may be monitored include server delay, packet counts and data rates, and frame relay FECN/BECN (Forward/Backward Explicit Congestion Notification) counts. Also "policy-related" metrics, such as, for example, how often "guaranteed rate" bandwidth failed to be allocated because there was insufficient free bandwidth, may be monitored.

In one embodiment, network device 120 monitors the customer network delay as well as the provider network delay to enable of service-

level management. Network device 120 monitors these delays as half round trip delays. In one embodiment, network device 120 monitors the inbound and outbound customer network delay, the inbound and outbound provider network delay, and host latency, and computes a single number which is the
5 sum of these three components. Note that the two customer network delay components do not measure the same request data packet that the inbound provider network delay does. They are computed from seeing the response data packet from the server.

In one embodiment, network device 120 includes a measurement
10 engine to record and maintain statistics and an on-board database to store information indicative of delays occurring in the network. The measurement engine may perform measurements and record information on site specific network delays, server delays, bits per second, lost packets, retransmission count, end-to-end delays, etc. Using this information, a
15 congestion index may be generated. The measurement engine may comprise hardware, software or a combination of both.

Network device 120 may also include management control that accesses the on-board database to determine whether a performance problem exists in the network that is the responsibility of the application

service provider. The management control may identify such a problem because a measurement taken by the measurement engine exceeds a predetermined value. In one embodiment, the management control may create histograms to provide a graphical representation of delays (e.g.,
5 average delays, cumulative delays, etc.).

If a problem is determined to exist, then the management control may notify the application service provider. In one embodiment, in response to the notification, the application service provider may send an event to notify the customer of the problem and/or offer a service that fixes or alleviates the
10 problem. In an alternative embodiment, the management control sends the event to the customer.

Once a problem attributed to the application service provider's infrastructure has been identified, the application service provider may remedy the situation. For example, if a particular class of traffic is causing
15 congestion, then more bandwidth can be allocated to prioritize traffic differently. Alternatively, network device 120 may shape the traffic to rectify the situation by, for example, controlling competing, non-essential, traffic. That is, traffic may be quantized depending on how much

bandwidth a particular user requires for the needs of the various applications.

In an alternative embodiment, network device 120 is installed between a service provider and an affiliate partner, such as, for example, a hosting service provider, application service provider, or a network service provider. In such a case, network device 120 performs mediation between the two providers. In still another alternate embodiment, network device 220, or the point of demarcation, is located at the point of interconnection between administrators.

Figure 2 illustrates an exemplary network environment in which an application demarcation point is utilized. Referring to Figure 2, an application service provider (ASP) utilizing a data center 202 is on one side of application demarcation point 201, while a customer data center 203 is on the opposite side of application demarcation point 201. Customer data center 203 comprises a local area network (LAN)(241) coupling computer systems 242_{1-N} and a network device 220. In one embodiment, network device 220 at customer data center 203 couples ASP data center 202 to a customer data center local area network (LAN) 241.

In one embodiment, the customer data center LAN 241 is also coupled to a customer branch 210 via a wide area network (WAN) 260. Customer branch 210 may comprises one or more computer systems networked together.

- 5 Network device 220 gathers information to ensure that the service is delivered to the customer (e.g., customer branch 210). In one embodiment, network device 220 generates a measurement that quantifies network congestion on either side of application demarcation point 201. This allows an application service provider, via a network manager, to isolate a problem
- 10 to one side of the network device. In one embodiment, this is performed in a non-intrusive way so that traffic flow is not interrupted.

In one embodiment, network device 220 may also use the information to create information to convince a customer that the service desired from the service provider has been delivered.

- 15 Figure 3 is an alternative view of Figure 1. Referring to Figure 3, network device 120 monitors the TCP flows that traverse the network between client 151 on the customer side of the network and server 150 on the provider side of the network. With TCP traffic, there are a series of data packets flowing in both directions. Such traffic flows are common,

particularly with respect to web-based applications. The flows are such that data packets travel in one direction while acknowledgement packets flow in the other direction. Some acknowledgement packets may include data that is traveling in the other direction and vice versa.

- 5 Each of the data packets have sequence numbers, which are indications of how many bytes have been transferred since the beginning of the connection where an "initial sequence number" was negotiated (which is not usually 0). The sequence numbers are usually in increasing order, except where there has been a retransmission or packets are arriving out-of-order.
- 10 The sequence numbers used in Figure 3 are not actual sequence numbers. Instead, for purposes of explanation, a simplified set of numbers have been used.

- Figure 3 shows data packets 1, 2, 3, and 4 traveling through network device 120 to server 150 (in that order) with associated acknowledgment
- 15 packets 1 and 3 traveling from server 150 towards network device 120. (Note that there are no acknowledgement packets for packets 2 and 4 as acknowledgement packets may not be generated for all data packets in a TCP-based network.) Also, packets 5, 7 and 8 travel from client 151 towards

network device 120. There is no packet 6 to indicate that packet 6 has been lost.

In one embodiment, network device 120, being at an application demarcation point, records the time it encounters packets and their acknowledgements to generate a congestion index. The congestion index is a measurement that quantifies network congestion of either side of the demarcation point. One embodiment of the process used by network device 120 to generate the congestion index is given below.

10 Exemplary Network Device Software

Figure 4 illustrates one embodiment of software code executed by one embodiment of network device 120. It should be noted that the operation performed by the software could be implemented in hardware (e.g., logic, circuitry, etc.) as well.

15 Referring to Figure 4, starting at line 3, there are a number of declarations to define variables. Starting at line 31, a structure is defined to keep track of sequence numbers and time stamps on data flows traveling in one direction. Each TCP flow is comprised of two half-flows going in two

directions around the same end point with data, at times, being piggy backed onto acknowledgement packets.

At line 60, a test determines if the packet has data in it. If it doesn't have data, processing transitions to line 85.

- 5 At lines 61-62, a data structure is allocated to hold sequence numbers and their associated time stamps, if such a data structure has not already been allocated. That is, the first time a packet with data is encountered on a half-flow, a data structure is created to store sequence numbers and their time stamps. Two data structures are allocated for a TCP flow with the dir
- 10 subcomponent indicating the direction of the flow (e.g., inbound, outbound).

Starting at line 65, the packet sequence number is added to the number of bytes of data to compute the sequence number of the last byte of data in the packet, which is the number that will be contained in the matching acknowledgement packet.

- 15 Starting at line 69, the comparison is made between the sequence number of the current packet and the last sequence number that has been seen. If the current sequence number is less than the last sequence number, then it is a duplicate and it is ignored.

Thus, at line 75-76, the sequence number of the packet and the time the packet reached the demarcation point are recorded in a data structure. In one embodiment, this data structure is a fixed size. In one embodiment, a packet depth of 7 is used, which is sufficient to get a good sample of packets from most connections and where a size greater would require at least twice the memory size. Even so, other sizes may be used.

If the packet is not a new packet, the network device processing transitions from line 69 to line 85. Starting at line 85, the packet is examined to see if its acknowledgement flag is set by the sender of the packet. If the packet is not an acknowledgement, the process is done.

If the packet contains an acknowledgement and a data structure is allocated (because data packet times have been recorded)(line 87), the acknowledgement is processed (lines 88-118); otherwise, processing on the packet is finished.

Starting at line 88, the software checks to see if the sequence number of the acknowledgement is not greater than the last acknowledgement that was seen. This would imply that the acknowledgement is a duplicate and should be ignored. Starting at line 91, the process loops through a number of locations in the data structure recording sequence numbers to find if a

time has been remembered for the data packet this packet is acknowledging.

At line 94, if the sequence number for the acknowledgement is equal to the sequence number of one of the previously recorded packets and the acknowledgement is a naked acknowledgment (i.e., the packet does not

5 include piggyback data), then the software performs a measurement.

Starting at line 96, the software calculates the difference between the time to see the acknowledgement at the demarcation and the time to initially see the data packet. This measurement is converted into milliseconds and refers to the time it took a packet to get from the demarcation point to the

10 other end and the matching acknowledgement to come back. This is the congestion index, which in this embodiment is computed as a millisecond delay. In summary, this method computes a number for a TCP half-flow for each different traffic class. Each half-flow has already been given a classification (e.g., traffic class by the classification engine described below),

15 which is retrieved at line 98. The software stores the congestion index measurement in a database (e.g., database 505 described below) associated with this traffic class.

Once a sample is recorded, starting at line 106, sequence number location in the data structure is set to zero (cleared) so that it can be used for

another sequence number of a packet that traverses the network. An optimization at line 115 causes the data structure to be completely cleared when there are no longer samples to process, because there is no outstanding data and the network device is not waiting on an
5 acknowledgement.

Block diagram of Device

Figure 5 is a data flow diagram of one embodiment of a network device described herein. Referring to Figure 5, a classification engine 501
10 classifies traffic in the traffic flow. Classification may be performed by using a classification tree. The traffic may be classified by all types of metrics. A classification model allows measurements of one type of traffic versus another type of traffic, one application versus another application, one customer versus another customer.

15 After classification, a response time block 502 monitors response time. Next, the traffic goes through shaping block 503 that performs any necessary shaping on the traffic. For more information on shaping, see U.S. application serial number 08/977,376, entitled "Method for Managing Flow Bandwidth Utilization at Network, Transport and Application Layers," filed

_____, incorporated herein by reference and assigned to the corporate assignee of the present invention.

Once classified, measurement engine 504 performs measurements to make determinations on comparable amounts of bandwidth. Essentially, the data is gathered, classified (e.g., on a per class basis), and then measurements are taken and grouped with other measurements associated with the same classification (e.g., according to application, application subtype, customer, subnet, etc.).

Measurement engine 504 provides probes to each of these blocks and stores measurements in embedded database 505. Measurement engine 504 records the congestion index sample in a “cumulative” variable as well as recording the number of samples in the bin corresponding to the traffic class. Each of the bins relates to a different classification type (e.g., application type). From these two measurements, an average congestion index can be determined.

Management control 505 fetches values from the embedded database 504 and quantifies and qualifies performance of the network based on these values. The results of the performance may be displayed through user interface 506.

For instance, a Citrix class of traffic may be monitored traffic in the inbound direction is measured as well as acknowledgements going in the other direction. The numbers recorded in the outbound Citrix flow reflect the delay or congestion index on one half of the traffic flow while the number recorded on the inbound reflect the delay on the other half. Thus, the Citrix flows are aggregated into 2 numbers, one on the inbound side (e.g., client portion 111) and one on the outbound side (e.g., provider portion 110). A comparison of the two may indicate whether a problem exists that is the responsibility of an application service provider.

The congestion index or other measurement values may be used to diagnose problems in the network and a determination of whether the problem is on the network provider side of the network or the customer side of the network. The diagnosis of a problem may be based on the fact that there is a difference in the congestion index values (or other measurement values) over time, based on the fact that the ratio of the congestion index (or other measurement values) on the provider side versus the same on the customer side are different, or whether such a ratio has changed, based on the fact that the congestion index (or other measurement values) on both sides has changed by the same amount (so as to diagnose that the problem

may be one of heavy traffic as opposed to a problem in infrastructure). A diagnosis may also be based on detected variances in the congestion index (or other measure(s)) over time for the same network time period or variances in the congestion index (or other measure(s)) between different

5 types of traffic.

Whereas many alterations and modifications of the present invention will no doubt become apparent to a person of ordinary skill in the art after having read the foregoing description, it is to be understood that any particular embodiment shown and described by way of illustration is in no

10 way intended to be considered limiting. Therefore, references to details of various embodiments are not intended to limit the scope of the claims which in themselves recite only those features regarded as essential to the invention.

CLAIMS

We claim:

- 1 1. A method comprising:
2 monitoring end-to-end performance of a network application at an
3 application demarcation point in a network; and
4 mediating between provider infrastructure and customer
5 infrastructure based on results of monitoring.
- 1 2. The method defined in Claim 1 wherein mediating between
2 provider and customer infrastructures comprises mediating between
3 performance impacts of underlying service layers and service
4 boundaries.
- 1 3. The method defined in Claim 1 further comprising
2 developing service level agreements underlying with a service boundary.
- 1 4. The method defined in Claim 1 wherein monitoring end-to-
2 end performance of the network application comprises measuring
3 performance of the network application.

1 5. The method defined in Claim 4 further comprising
2 measuring latency, inbound and outbound provider network delay, and
3 inbound and outbound customer network delay.

1 6. The method defined in Claim 4 further comprising
2 diagnosing a problem based on the measurements.

1 7. The method defined in Claim 4 wherein measuring
2 performance of the network application comprises measuring both
3 quantitative performance of the network application.

1 8. The method defined in Claim 1 further comprising
2 determining from where a service effecting problem is emanating in the
3 network.

1 9. The method defined in Claim 1 further comprising a service
2 provider creating, monitoring, and managing service level agreements
3 for a delivered end-to-end service delivered which traverses its
4 infrastructure in the network.

ABSTRACT OF THE DISCLOSURE

- A method and apparatus for using an application layer demarcation point are described. In one embodiment, the method comprises monitoring end-to-end performance of a network application at an application
- 5 demarcation point in a network, and mediating between provider infrastructure and customer infrastructure based on results of monitoring.

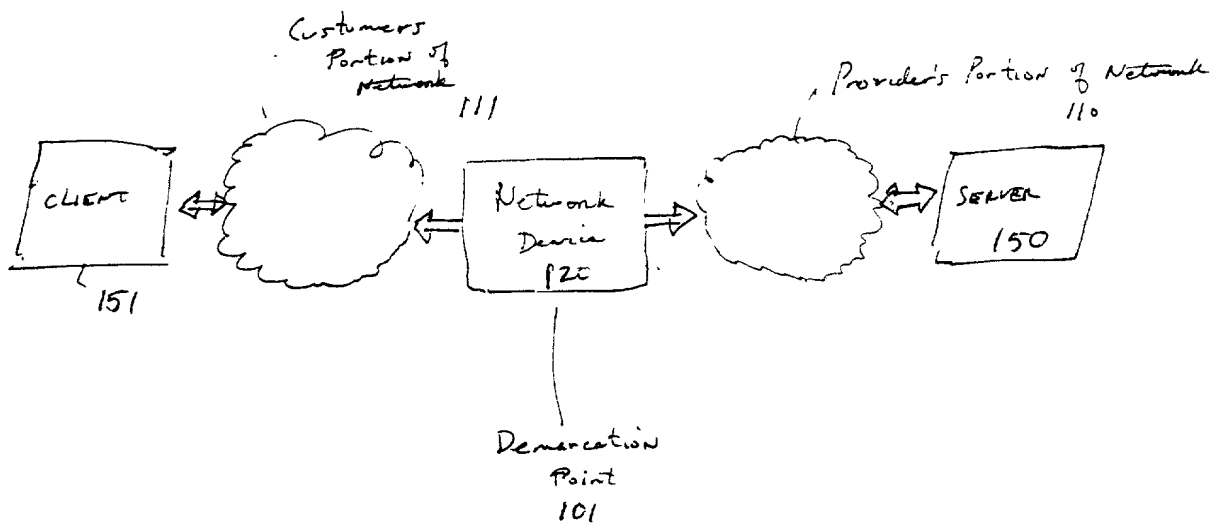
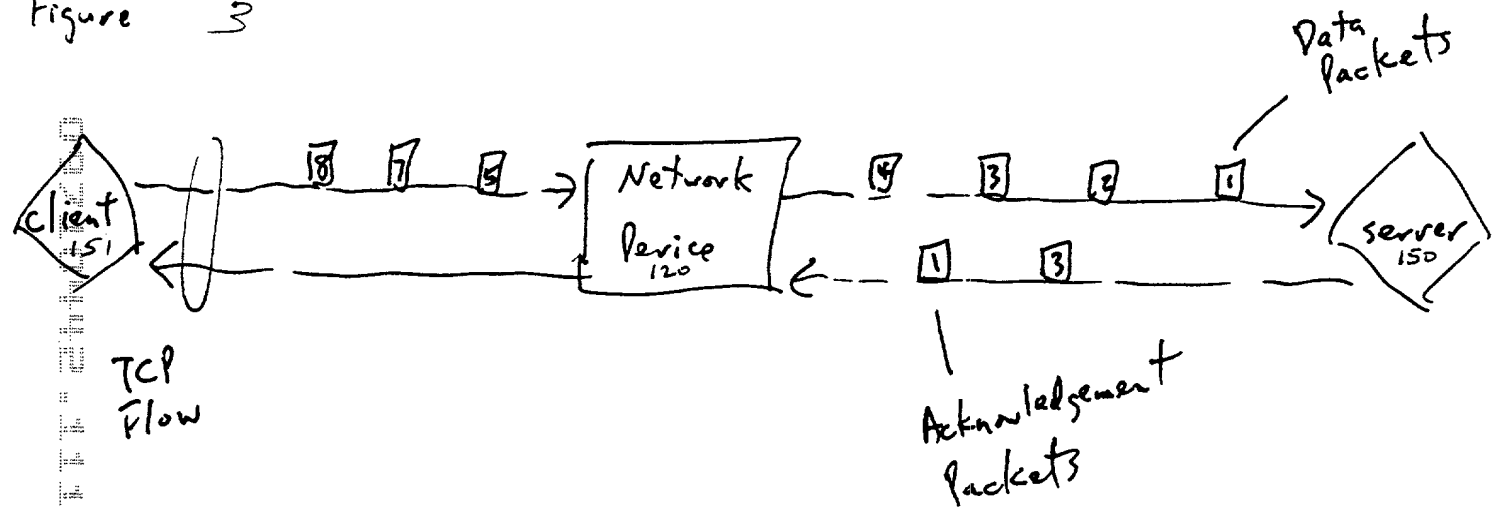


Figure 1

The diagram illustrates a network architecture and its performance metrics. The top part shows a network flow from a Customer WAN (260) through a Customer Data Center LAN (242) to an Application Demarc (201) and finally to a Network device (220). The bottom part shows a timeline of delays: Inbound Customer Network Delay, Inbound Provider Network Delay, Outbound Customer Network Delay, and Outbound Provider Network Delay, all contributing to the End-to-End Response Time - 'Quality of Experience' Correlation (210).

Figure 3



SCANNED, #
drixoral.c

drixoral.c

```

73 if(i < DRIX_SEQ_SPOTS){
74     drix->seq[i] = seq;
75     drix->time[i] = info->bcb->tick.ticks;
76     }
77     else
78     info2(*req &d spot overflow tcb &X', seq, tcb);
79     drix->lastseq = seq;
80     }
81     )
82     )
83     )
84     )
85     dir = OTHER_DIRECTION(info->dir);
86     if(IS_FLAG_SEQ(info->header.flags, TCP_FLAG_ACK) && (drix = am->drixoral[dir])
87     && (SEQ_OT(info->header.ack, drix->lastack) || (drix->lastack))){
88     leave = 0;
89     for(i = 0; i < DRIX_SEQ_SPOTS; i++)
90     if((drix->seq[i])
91     continue;
92     else if(SEQ_SEQ(info->header.ack, drix->seq[i])){
93     if(info->header.dataLen == 0){
94     msec = TICKS_TO_MSECS_ROUNDED(info->bcb->tick.ticks - drix->time[i]);
95     msec = TICKS_TO_MSECS_ROUNDED(info->bcb->tick.ticks - drix->time[i]);
96     }
97     if((tc = tcClassIdToClassFast(tcb->halfConn[dir].gear.classId))){
98     metClassCongestionAccum(tc, msec);
99     blurt4(*sample &d msec &d class &d dir &d", i, msec, tc->name, dir);
100     }
101     else
102     attnl("no class for TCB &X", tcb);
103     )
104     )
105     drix->seq[i] = 0;
106     }
107     else if(SEQ_OT(info->header.ack, drix->seq[i])){
108     drix->seq[i] = 0;
109     else
110     leave++;
111     }
112     if(leave)
113     drix->lastack = info->header.ack;
114     else{
115     kfree(drix);
116     am->drixoral[dir] = 0;
117     )
118     )
119     )
120     void
121     drixoralCleanup(
122     TCB_PTR tcb
123     ){
124     ArtMeasurementPtr am;
125     CongestionMeasurementPtr drix;
126     am = tcb->artData;
127     if((drix = am->drixoral[DIR_INBOUND])){
128     kfree(drix);
129     am->drixoral[DIR_INBOUND] = 0;
130     )
131     )
132     if((drix = am->drixoral[DIR_OUTBOUND])){
133     kfree(drix);
134     am->drixoral[DIR_OUTBOUND] = 0;
135     )
136     )
137     )
138     )
139     )
140     }
141     }

```

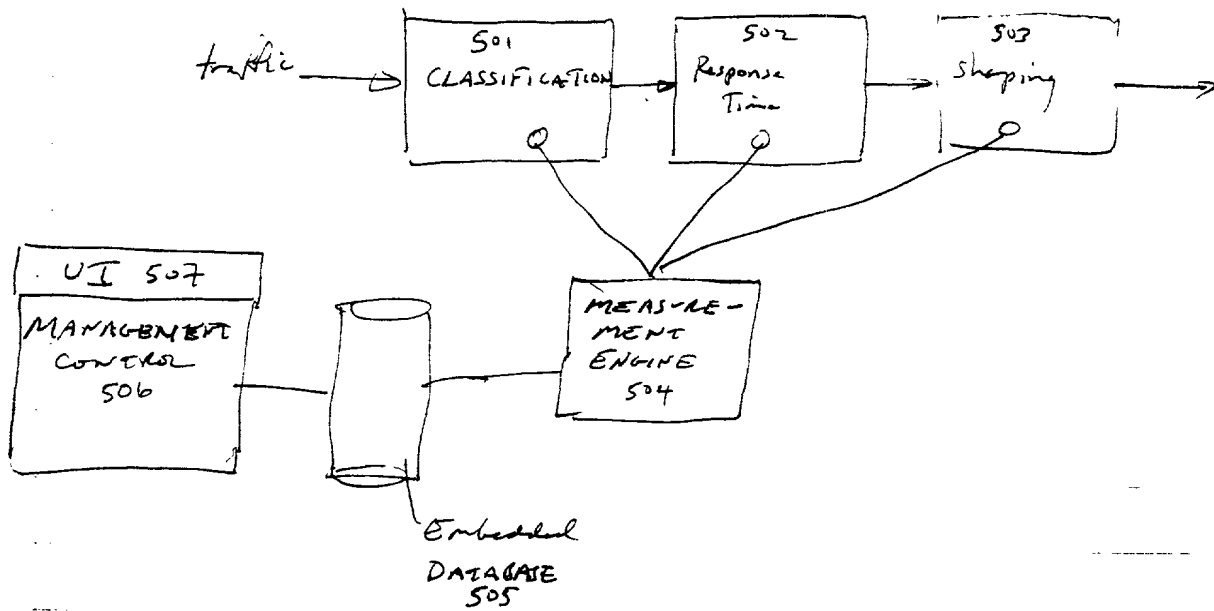


Figure 5

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below, next to my name.

I believe I am the original, first, and sole inventor (if only one name is listed below) or an original, first, and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

"APPLICATION SERVICE LEVEL MEDIATION AND METHOD OF USING THE SAME"

the specification of which

 X is attached hereto.
 was filed on (MM/DD/YYYY) _____ as
United States Application Number _____
or PCT International Application Number _____
and was amended on (MM/DD/YYYY) _____.
(if applicable)

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claim(s), as amended by any amendment referred to above. I do not know and do not believe that the claimed invention was ever known or used in the United States of America before my invention thereof, or patented or described in any printed publication in any country before my invention thereof or more than one year prior to this application, that the same was not in public use or on sale in the United States of America more than one year prior to this application, and that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representatives or assigns more than twelve months (for a utility patent application) or six months (for a design patent application) prior to this application.

I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d), of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

<u>Prior Foreign Application(s)</u>			<u>Priority Claimed</u>	
_____ (Number)	_____ (Country)	_____ (Foreign Filing Date - MM/DD/YYYY)	Yes	No
_____ (Number)	_____ (Country)	_____ (Foreign Filing Date - MM/DD/YYYY)	Yes	No
_____ (Number)	_____ (Country)	_____ (Foreign Filing Date - MM/DD/YYYY)	Yes	No

I hereby claim the benefit under title 35, United States Code, Section 119(e) of any United States provisional application(s) listed below:

_____ (Application Number)	_____ (Filing Date – MM/DD/YYYY)
_____ (Application Number)	_____ (Filing Date – MM/DD/YYYY)

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

_____ (Application Number)	_____ (Filing Date – MM/DD/YYYY)	_____ (Status -- patented, pending, abandoned)
_____ (Application Number)	_____ (Filing Date – MM/DD/YYYY)	_____ (Status -- patented, pending, abandoned)

I hereby appoint the persons listed on Appendix A hereto (which is incorporated by reference and a part of this document) as my respective patent attorneys and patent agents, with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

Send correspondence to Michael J. Mallie, BLAKELY, SOKOLOFF, TAYLOR &
(Name of Attorney or Agent)
ZAFMAN LLP, 12400 Wilshire Boulevard 7th Floor, Los Angeles, California 90025 and direct
telephone calls to Michael J. Mallie, (408) 720-8300.
(Name of Attorney or Agent)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Sole/First Inventor Todd Krautkremer

Inventor's Signature _____ Date _____

Residence _____ (City, State) _____ Citizenship _____ (Country) _____

Post Office Address _____

Full Name of Second/Joint Inventor Guy Riddle

Inventor's Signature _____ Date _____

Residence _____ (City, State) _____ Citizenship _____ (Country) _____

Post Office Address _____

Full Name of Third/Joint Inventor _____

Inventor's Signature _____ Date _____

Residence _____ (City, State) _____ Citizenship _____ (Country) _____

Post Office Address _____

Full Name of Fourth/Joint Inventor _____

Inventor's Signature _____ Date _____

Residence _____ (City, State) _____ Citizenship _____ (Country) _____

Post Office Address _____

APPENDIX A

William E. Alford, Reg. No. 37,764; Farzad E. Amini, Reg. No. 42,261; William Thomas Babbitt, Reg. No. 39,591; Carol F. Barry, Reg. No. 41,600; Jordan Michael Becker, Reg. No. 39,602; Lisa N. Benado, Reg. No. 39,995; Bradley J. Berezna, Reg. No. 33,474; Michael A. Bernadico, Reg. No. 35,934; Roger W. Blakely, Jr., Reg. No. 25,831; R. Alan Burnett, Reg. No. 46,149; Gregory D. Caldwell, Reg. No. 39,926; Andrew C. Chen, Reg. No. 43,544; Thomas M. Coester, Reg. No. 39,637; Donna Jo Coningsby, Reg. No. 41,684; Florin Corie, Reg. No. 46,244; Dennis M. deGuzman, Reg. No. 41,702; Stephen M. De Klerk, Reg. No. 46,503; Michael Anthony DeSanctis, Reg. No. 39,957; Daniel M. De Vos, Reg. No. 37,813; Sanjeet Dutta, Reg. No. 46,145; Matthew C. Fagan, Reg. No. 37,542; Tarek N. Fahmi, Reg. No. 41,402; George Fountain, Reg. No. 37,374; James Y. Go, Reg. No. 40,621; James A. Henry, Reg. No. 41,064; Libby N. Ho, Reg. No. 46,774; Willmore F. Holbrow III, Reg. No. 41,845; Sheryl Sue Holloway, Reg. No. 37,850; George W. Hoover II, Reg. No. 32,992; Eric S. Hyman, Reg. No. 30,139; William W. Kidd, Reg. No. 31,772; Sang Hui Kim, Reg. No. 40,450; Walter T. Kim, Reg. No. 42,731; Eric T. King, Reg. No. 44,188; George Brian Leavell, Reg. No. 45,436; Kurt P. Leyendecker, Reg. No. 42,799; Gordon R. Lindeen III, Reg. No. 33,192; Jan Carol Little, Reg. No. 41,181; Robert G. Litts, Reg. No. 46,876; Joseph Lutz, Reg. No. 43,765; Michael J. Mallie, Reg. No. 36,591; Andre L. Marais, under 37 C.F.R. § 10.9(b); Paul A. Mendonsa, Reg. No. 42,879; Clive D. Menezes, Reg. No. 45,493; Chun M. Ng, Reg. No. 36,878; Thien T. Nguyen, Reg. No. 43,835; Thinh V. Nguyen, Reg. No. 42,034; Dennis A. Nicholls, Reg. No. 42,036; Robert B. O'Rourke, Reg. No. 46,972; Daniel E. Ovanezian, Reg. No. 41,236; Kenneth B. Paley, Reg. No. 38,989; Gregg A. Peacock, Reg. No. 45,001; Marina Portnova, Reg. No. 45,750; William F. Ryann, Reg. No. 44,313; James H. Salter, Reg. No. 35,668; William W. Schaal, Reg. No. 39,018; James C. Scheller, Reg. No. 31,195; Jeffrey Sam Smith, Reg. No. 39,377; Maria McCormack Sobrino, Reg. No. 31,639; Stanley W. Sokoloff, Reg. No. 25,128; Judith A. Szepesi, Reg. No. 39,393; Vincent P. Tassinari, Reg. No. 42,179; Edwin H. Taylor, Reg. No. 25,129; John F. Travis, Reg. No. 43,203; Joseph A. Twarowski, Reg. No. 42,191; Tom Van Zandt, Reg. No. 43,219; Lester J. Vincent, Reg. No. 31,460; Glenn E. Von Tersch, Reg. No. 41,364; John Patrick Ward, Reg. No. 40,216; Mark L. Watson, Reg. No. 46,322; Thomas C. Webster, Reg. No. 46,154; and Norman Zafman, Reg. No. 26,250; my patent attorneys, and Firasat Ali, Reg. No. 45,715; Justin M. Dillon, Reg. No. 42,486; Thomas S. Ferrill, Reg. No. 42,532; and Raul Martinez, Reg. No. 46,904, my patent agents, of BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, with offices located at 12400 Wilshire Boulevard, 7th Floor, Los Angeles, California 90025, telephone (310) 207-3800, and James R. Thein, Reg. No. 31,710, my patent attorney with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

APPENDIX B

Title 37, Code of Federal Regulations, Section 1.56 Duty to Disclose Information Material to Patentability

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclose information exists with respect to each pending claim until the claim is cancelled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is cancelled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclose all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

(1) Prior art cited in search reports of a foreign patent office in a counterpart application, and

(2) The closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made or record in the application, and

(1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or

(2) It refutes, or is inconsistent with, a position the applicant takes in:

(i) Opposing an argument of unpatentability relied on by the Office, or

(ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

(1) Each inventor named in the application;

(2) Each attorney or agent who prepares or prosecutes the application; and

(3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.

(d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.